

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY (Leave Blank)	2. REPORT DATE 5/29/98	3. REPORT TYPE AND DATES COVERED Final Report 6/1/95 - 5/31/98		
4. TITLE AND SUBTITLE Generation of Tutorial Dialogues: Discourse Strategies for active Learning		5. FUNDING NUMBERS N00014-95-1-1078		
6. AUTHORS Dr. Martha Evens				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Illinois Institute of Technology 3300 S. Federal Street, Room 301 Chicago, IL 60613		8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Office of Naval Research ONR 251: C, Ballston Tower One 800 North Quincy Street, Arlington, VA 22217-5660		10. SPONSORING / MONITORING AGENCY REPORT NUMBER		
11. SUPPLEMENTARY NOTES		19980605 116		
12a. DISTRIBUTION / AVAILABILITY STATEMENT Unlimited		12b. DISTRIBUTION CODE		
<div style="border: 1px solid black; padding: 5px; text-align: center;"> DISTRIBUTION STATEMENT A Approved for public release; Distribution Unlimited </div>				
13. ABSTRACT (Maximum 200 words) <p>With the support of the Cognitive Science Program of ONR, we are developing the capability to generate complex natural language tutorial dialogues for an intellignet tutoring system designed to help medical students understand the functioning of the negative feedback system that regulates blood pressure in the human body. We are convinced that a real natural language interface is vital to a tutoring system trying to help students learn complex concepts like negative feedback. The text generator for a tutoring system must be able to ask questions and provide hints in addition to generating definitions, descriptions, and explanations of the functioning of the physiological system and the underlying anatomy. The language understanding component must be ready to accept student responses and student initiatives, which are full of ellipses, novel abbreviations, wild spellings, and wilder grammar.</p> <p>Our work is embodied in a system called Circsim-Tutor. Briefly described, Circsim-Tutor presents the student with a set of clinical problems each of which results in a perturbation of blood pressure, and asks the student to explain step-by-step how the blood pressure is perturbed and how the perturbation is physiologically compensated for. The system conducts a tutorial dialogue in English, as the student does this, with the session organized around the students's errors in making these predictions.</p>				
14. SUBJECT TERMS		15. NUMBER OF PAGES 12		
		16. PRICE CODE		
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT Unclassified	

REPORT ON AASERT SUPPLEMENTS TO GRANT No. N00014-89-J-1952,
GRANT AUTHORITY IDENTIFICATION NUMBER NR4422554 (1992-1995),
AND GRANT No. N00014-94-1-0338 (1995-1998)
FROM THE COGNITIVE SCIENCE PROGRAM, OFFICE OF NAVAL RESEARCH
TO ILLINOIS INSTITUTE OF TECHNOLOGY.

Martha Evens
Computer Science, Illinois Institute of Technology
10 West 31st Street, Room 236, Chicago, IL 60616

With the support of the Cognitive Science Program of the Office of Naval Research, we are developing the capability to generate complex natural language tutorial dialogues for an intelligent tutoring system designed to help medical students understand the functioning of the negative feedback system that regulates blood pressure in the human body. We are convinced that a real natural language interface is vital to a tutoring system trying to help students learn complex concepts like negative feedback in two ways. First of all, teaching new concepts is inextricably intertwined with the teaching of the language in which those concepts are expressed; real understanding involves an ability to describe what a system is doing in appropriate language. Second, before the tutor can be certain that students really understand the material, that they are not just successfully playing a video game, they must describe their knowledge in language. In order to provide these capabilities we must develop significant student modeling facilities as well as components for natural generation and understanding, that allow the machine tutor to understand and comment on self-explanations from the student.

The text generator for a tutoring system must be able to ask questions and provide hints in addition to generating definitions, descriptions, and explanations of the functioning of the physiological system and the underlying anatomy. The language understanding component must be ready to accept student responses and student initiatives, which are full of ellipses, novel abbreviations, wild spelling, and wilder grammar. We have found that Lexical Functional Grammar as developed by Bresnan and Kaplan to provide a helpful theoretical approach to these tasks. Extensive sublanguage analysis of actual human keyboard-to-keyboard tutoring sessions has provided the basis for the lexicon and the grammar.

Our work is embodied in a system called Circsim-Tutor. Briefly described, Circsim-Tutor presents the student with a set of clinical problems each of which results in a perturbation of blood pressure, and asks the student to explain step-by-step how the blood pressure is perturbed and how the perturbation is physiologically compensated for. The student is presented with a table of seven physiologically significant variables, for which the student is to predict qualitatively, first, the direct physical effects of the perturbation, reflecting how these seven variables affect one another, second, how the body uses the negative feedback system controlled by the baroreceptor reflex to compensate for these changes, and, third, what the ultimate steady state is for each. The system conducts a tutorial dialogue in English, as the student does this, with the session organized around the student's errors in making these predictions.

When we first applied for support from the AASERT program we had just managed to put together the first working version of our system complete with natural language modules. We were generating sentences, but we were just beginning to understand dialogue issues above the sentence level, and we were motivated by the need for help with these issues. It was clear, also, that we had just begun to understand the vast amount of information in the tutoring dialogues that were being produced by Joel Michael and Allen Rovick.

Gregory Sanders had just finished the course work for the Ph.D. in June, 1992, when the first AASERT funding began. He had decided that he wanted to work on natural language processing and was talking to Peter Greene as well as to me about topics. He was also teaching Systems Programming at IIT part-time.

Sanders started out looking for multiturn structures in the human tutoring sessions. He was the first of us to discover "Directed Lines of Reasoning," or DLRs, explanations and summaries with the same content as monologues, but delivered interactively as a series of questions and answers. He developed schemas for summaries of the reasoning in the DR stage, in the RR stage, and in the SS stage. He also developed schemas for remediating explanations for several of the most common explanations. These schemas were designed so that the system could deliver them interactively or as monologues depending on the state of the student model.

Sanders was also the first to recognize and study another important discourse phenomenon, that we were not set up to handle: student initiatives, in which the student tries to change the course of the dialogue. He made additions to the instructional planner and the text generator to allow the system to respond to one of the most frequent types of initiatives, requests for definitions in the forms:

I am confused about <variable name>.

I do not understand <variable name>.

What is <variable name>?

And he persuaded Chung Hee Lee to add the necessary logic form to the Input Understander.

After he received his degree Greg Sanders taught at Hood College in Frederick, Maryland, for two years and then decided that he would rather do than teach. He found a job with NIST.

Gregory Hume also started working with us in the summer of 1992. At that point he was an instructor of computer science in the Mathematics department of Valparaiso University, looking for a Ph.D. topic in artificial intelligence. He started to read about student modeling and was fascinated by the work of Kurt Van Lehn. He was instrumental in arranging for Van Lehn's trip to Chicago. While Kurt was interviewing medical students during tutoring sessions, Hume interviewed Joel Michael. As Greg watched, Joel recognized that the student he was tutoring keyboard-to-keyboard style was confused and upset by hints and himself stopped issuing hints. This was the first time that we had realized what a conscious process hint construction is for Joel and for Allen Rovick. Further interviews with Michael and Rovick helped Hume to understand the wide scope of hints and their central importance in

the tutoring process. This led to a detailed study of hinting in the transcripts and a series of publications.

On the basis of these publications Hume persuaded Valparaiso to give him a tenure-track position, and this year he obtained tenure. His next goal is an independent computer science department.

Reva Freedman arrived to join us just as Greg Sanders and Greg Hume were finishing their theses. She realized that we needed to redesign the instructional planner and the discourse planner to function together if we were going to use the multiturn schemas designed by Sanders and deliver the hinting strategies discovered by Hume. Freedman redefined the whole instructional planner in terms of hierarchical sets of rules. She wrote sets of rules for tutoring strategies and for tutoring tactics to carry out those strategies. She reorganized a large part of the knowledge base in terms of rules as well.

Freedman implemented the planner rules and the discourse generation rules as Longbow operators. We have recently received a new version of Longbow from Michael Young and Johanna Moore. This planner is much more useful for our purposes than UCPOP and Johanna Moore's generosity in sharing it with us from its inception has been extremely helpful.

Freedman rewrote the Problem Solver in terms of planning operators as well. This should make it much easier for Joel Michael and Allen Rovick to add new procedures. We look on this as the first step toward an authoring tool.

Freedman also persuaded us to start marking up the tutoring transcripts in SGML style. This has made it possible for us to produce actual counts of alternative strategies used by the tutors and to use machine learning programs.

Freedman started her work on a Ph.D. with Gilbert Krulee at Northwestern University, but they had not been able to agree on a topic. She reached the time limit at Northwestern and applied to IIT. On the basis of her work with us she was eventually reinstated at Northwestern and received her Ph.D. from there.

Stefan Brandle began by building a new interface for version 3, which he then retrofitted to version 2 in both the Macintosh and PC versions. The most important feature of this new interface was the interleaving of student's and the tutor's portions of the dialogue. Retrofitting this new interface to Version 2 made a large qualitative difference to users of the system - it seemed to be a new system. All the windows are scrollable so that users can review the instructions, reread the perturbation in question, and go over the past conversation. He redid the fonts as well and made the screen much more readable.

Brandle was also horrified by Circsim-Tutor's acknowledgments. The original Version 2 responded to every student answer with "Correct" or "Wrong." Brandle determined to rectify this situation. He started with a temporary *ad hoc* fix: He added some alternatives

for the use of Kumar Ramachandran's lexical choice program, but he determined find a theoretical basis for a study of the ways in which Michael and Rovick use acknowledgments.

He began by looking at acknowledgement strategies in telecommunications protocols and how they establish and defer closure. He found a number of parallels between the two situations.

Then Brandle heard Herbert Clark talk at the Discourse Processes meeting in San Diego in 1996 and found the theoretical basis he needed in Clark's theory of joint actions. He decided to define acknowledgments in the light of Herbert Clark's presentation of discourse as a joint activity composed of joint actions. Clark focuses on the mutual coordination of their individual actions by the participants in linguistic activity. In tutoring sessions as in telecommunication sessions a signal of closure at one level implies that closure has succeeded at all lower levels. As Clark points out, it is generally not possible to succeed at the level of uptake on the joint action (Level 4) without having understood the signal (Level 3), without having decoded the signal (Level 2), without having paid attention to the signal (Level 1).

Brandle carried out three experiments in which we asked several informants to make judgments about the role of acknowledgments in providing or deferring closure and about the ways in which closure was signalled in human tutoring sessions. He obtained fairly good inter-rater reliability with kappa values between seventy and eighty percent for markedness and closure alternatives and higher than that for judgments of whether a particular acknowledgment is positive, negative, or neutral.

Brandle had taught part-time both at Wheaton College and at IIT before he started to work with us. He taught at IIT full-time for a year with marvellous evaluations and decided to resign to work on his thesis. He is now trying to choose between assistant professorships at IIT and at Wheaton.

Michael Glass resigned from a full-time job at the Fermi National Accelerator Laboratory to work with us on the Circsim-Tutor project. He had already attained a first-class reputation and produced a number of publications in scientific computing. He had been teaching part-time at IIT for several years. When he came he also took over the whole gamut of problems involved in running the laboratory.

Glass identified a number of problems with our old approach of trying to interpret every student input as an answer to the question just posed by the system. First of all, this approach causes the system to reject "near misses," when the student gives an answer that is correct but unexpected. Suppose the system asks for the mechanism that controls TPR (total peripheral resistance), expecting the answer "nervous system" and the student answers "Arteriolar Resistance." This answer is not wrong, and if we reject it, we are likely to confuse students hopelessly. When this happens in the expert tutoring sessions, the tutor responds by asking what controls Ra.

Glass also analyzed the ways that students mix bits of equations with ordinary sentences. So we see:

CO is HR times SV

CO = HR times SV

CO equals hr*sv

But this seems to contradict $co = hr \times sv$, so

Glass has written an equation grammar and integrated with the sublanguage grammar used by the parser.

Glass also noticed that certain tutor hints have "evoking strength;" they tend to pull right answers out of students. The concepts "preload" and "afterload" seem to be particularly useful in this fashion. He proposed that we move these concepts to the middle level of the concept map and use them in understanding and generation both.

He has now embodied these ideas in a new input understander based on the information extraction techniques pioneered by the DARPA *Message Understanding Conferences*. This input understander has recently been used in a trial of Circsim-Tutor with twenty-two first year medical students at Rush Medical College. The old input understander responded with "I'm sorry but I did not understand you. Please rephrase." about ten percent of the time. Glass's new input understander was stumped less than one percent of the time in these trials, without any increase in the number of misunderstandings.

Titles and Abstracts of Ph.D. Theses Written by AASERT Students

Gregory Hume. *Using Student Modeling to Determine How and When to Hint in an Intelligent Tutoring System.* May, 1995.

I set out to design a student model for Circsim-TUTOR that accurately reflects the way a human tutor perceives a student's progress. While this model will build on past work in student modeling and user modeling, it will also reflect our analysis of the behavior of expert tutors as shown in the transcripts of human tutoring sessions collected by our collaborators, Joel Michael and Allen Rovick, who are Professors of Physiology at Rush Medical College.

Past research in student modeling has typically used an overlay model for representing declarative knowledge and a bug library for representing procedural knowledge. We need to combine these two concepts. The new knowledge base for Circsim-TUTOR, which I developed as a Ph.D. Qualifying Project, represents the problem solver in "runnable form" in the knowledge base, which makes it possible to overlay procedural knowledge as well as declarative knowledge. I have extended the concept of the bug library to cover both declarative and procedural misconceptions. Michael and Rovick seem to treat these misconceptions in the same way, whether they are declarative or procedural in nature: they identify them by certain patterns of responses; they classify them as fundamental (overriding the current tutoring goal) or derived (to be marked for attention later); they have certain tested methods of eradicating them or trying to; they expect students to have trouble changing these misconceptions and often check later to see whether the bug-oriented tutoring has been successful.

Study of human tutoring sessions has revealed that Michael and Rovick allow the student to make a number of predictions before they engage in an interactive dialogue. This allows them to identify key patterns of errors. These human tutors use a very coarse-grained student modelling scheme to evaluate the student. They seem to compute a global assessment that represent how well the student is doing overall and a local assessment that is recomputed every time the student starts in on a new topic.

Michael and Rovick constantly attempt to promote active learning. They regularly use hints and only resort to giving an explanation when they believe that the student will not be able to understand a hint. They rely heavily on the local assessment to decide when and how to hint. While the hints come in many surface forms, there are two significantly different categories of hints: convey- information hints (CI-hints) and point to information hints (PI-hints). Hinting is a very knowledge-intensive process; it requires extensions to our knowledge base. Hinting is a conscious process and the tutors stop hinting when a student has trouble responding to hints.

Gregory Sanders. *Generation of Explanations and Multi-Turn Discourse Structures in Tutorial Dialogue Based on Transcript Analysis.* July, 1995.

The goal of this research was to understand how human tutors generate sophisticated multiturn discourse and to propose schemas for an intelligent tutoring system to generate this kind of discourse. It seemed clear that research on tutoring discourse could substantially increase the effectiveness of Circsim-Tutor. The old version of Circsim-Tutor understood and generated English one sentence at a time, and the tutoring dialogue that resulted is substantially simpler and less effective than it could be. It is clear that human tutors use significantly more complex multiturn structures. Further, human tutors generate explanations and summaries that are several sentences in length.

We have obtained transcripts of approximately 32 hours of keyboard-to-keyboard tutoring sessions, and several hours of face-to-face sessions, with medical students from Rush Medical College being tutored by two Physiology Professors at Rush Medical College with extensive tutoring experience. As we look at multi-turn discourse units we discover many dialogue-level problems that we did not recognize earlier. How do our expert tutors respond to student initiatives? How does the tutor sustain a topic over several turns? When does the tutor summarize the preceding discussion? Study of our human tutors in action shows that they often deliver explanations and summaries interactively as Directed Lines of Reasoning, especially when the student is doing well. The same summary or explanation content is delivered as a monologue, when the student is doing badly.

Work done by Barbara Fox shows that repair, rather generally construed, is a crucial strategy in tutoring sessions. Our sessions with human tutors appear to support this. Can we figure out how to make a machine tutor carry out repair? It seems to require even more intelligence than the rest of the session. Here we have chosen an alternative - combining Fox' work and analysis of our own data to determine how to minimize repair situations instead.

We have carried out a detailed analysis and classification of student initiatives in the keyboard-to-keyboard tutoring sessions. In order to apply this analysis we need to make Circsim-Tutor recognize the student initiatives and their intent. Then we need to understand how our expert tutors craft the long interchanges that typically result from student initiatives in our data.

Reva Freedman. *Interaction of Discourse Planning, Instructional Planning, and Dialogue Management in an Interactive Tutoring System.* Dr. Freedman did her research with us but completed her thesis at Northwestern University with Gilbert Krulee. She defended her thesis in August, 1996.

We demonstrate the utility of natural language generation as the underlying model for an intelligent tutoring system (ITS) in cardiovascular physiology. We have achieved this goal by dividing it into three subgoals, each of which builds on its predecessor: (a) developing a model of the tutorial dialogue of human tutors based on current research in natural language

generation, with emphasis on text planning and the Conversation Analysis school, (b) analyzing a corpus of human-to-human tutoring sessions in cardiovascular physiology in terms of this model, and (c) designing an ITS that implements the model. We develop an abstract model of tutorial dialogue in order to put text generation for ITSs on solid theoretical footing. We give a detailed analysis of our corpus using this model, including a discussion of how tutors sequence their corrections, begin and end phases of the discourse, acknowledge responses, reply to student errors, teach different kinds of information, provide hints, conduct interactive explanations and choose between domain models. We present a detailed design for an ITS which uses this model to show that it can be implemented with current technology. The system is divided into two routines running in parallel, a global tutorial planner, which makes discourse decisions for units larger than a turn, and a turn planner, which assembles individual turns. The tutorial planner does not generate text directly, but generates a series of semantic forms. The turn planner collects the semantic forms for a turn, which may include information from multiple tutorial and/or conversational goals, and generates text for them as a unit. This architecture promotes coherent dialogue while permitting the tutor to use multi-turn discourse plans and change plans in response to student input. We expect this model to produce longer, more complex, and more varied dialogues than previous work.

Stefan Brandle, *Using Joint Actions to Explain Acknowledgments in Tutorial Discourse: Application to Intelligent Tutoring Systems*. May, 1998.

Human communication is powerful and successful, not entirely because of accuracy in understanding utterances, but also due to the ability of the communicating parties to dynamically detect and correct problems as they arise. Acknowledgments and associated communication mechanisms are key to this process of coordinating discourse. This research covers a small, but significant, aspect of using natural language for human-computer interaction (HCI). It deals with the problem of how humans coordinate their communication with other humans and with computers, through the use of acknowledgments and related phenomena. The primary goals are 1) to advance the understanding of coordination in communication by drawing together information from a number of different fields, 2) to describe coordination in tutorial dialogue, and 3) to propose a model of how intelligent tutoring systems (ITSs) should generate linguistic acknowledgments. The document presents a brief overview of intelligent tutoring systems. It then studies different approaches to building a theoretical foundation for understanding acknowledgments. In particular, it proposes that the joint action theory proposed by Herbert Clark forms a good foundation for a theory of acknowledgments. Clark's framework is further supported by ideas from linguistics and discourse theory, information theory, data communication, and a brief mention of verbal behavior analysis from the behaviorist tradition. Three initial experiments in identifying and categorizing acknowledgments are also discussed. A subset of these results is used in a machine-learning experiment to build a computational model for the generation of acknowledgments that can be used by intelligent tutoring systems such as CIRCSIM-Tutor, the system built by our research group. Last, there is a presentation of how acknowledgments and the idea of joint action is relevant to the development of human-computer interfaces.

Michael Glass. *Using Information Extraction Techniques to Understand Unconstrained Natural Language Input in an Intelligent Tutoring System (tentative).*

The new input understander described here is based on techniques developed for the DARPA *Message Understanding Conferences*. It is designed to handle unconstrained natural language input to our intelligent tutoring system, Cirsim-Tutor, which asks students to predict how the values of parameters controlling blood pressure will respond to a perturbation, and then carries on a remedial dialogue in natural language. These inputs are usually fragmentary and sometimes misspelled. Our old technique of combining the answer with the question to produce a semantic representation of the input failed utterly when the student gave an unexpected but correct answer or attempted to move the dialogues in a new direction. This new approach allows us to respond to unexpected responses of many kinds.

The grammar has been expanded to handle equations and algebraic and chemical expressions, which students often include in answers. The knowledge base has been expanded to support several types of unexpected but correct answers. The system also recognizes several kinds of student initiatives. Elmi's approach to automatic spelling correction and Dardaine's case frames have been adapted to support recognition of ill-formed input.

This input understander has been used in a test of Cirsim-Tutor with the whole body of first-year students in the alternative (problem-based) curriculum at Rush Medical College, where it was able to parse all but two student inputs. The old input understander responded to about one student input in ten with: "I am sorry but I do not understand you. Please rephrase." So the new system is much more comfortable to use. The new understander also gives a more detailed report of what the student intended, so that the system can generate more varied and more sophisticated answers.

Current Positions of AASERT Students

Gregory Sanders is a "Computer Scientist" at the National Institute of Standards and Technology (NIST). He is working in the Spoken Natural Language Processing Group (894.01) of the Information Technology Laboratory. Currently Greg and Dr. Jean Scholtz of NIST are engaged in creating an evaluation for Allen Sears' "DARPA Communicator" program. Jean Scholtz is trying to figure out how to measure and evaluate the interaction capabilities of implementations of Communicator, and Greg is trying to measure the effectiveness of the dialogue.

Gregory Hume is a tenured Associate Professor of Mathematics and Computer Science at Valparaiso University. He organized the Intelligent Tutoring Systems track of the FLAIRS Conference held last week. He is building a tutoring system that combines language and pictures.

Reva Freedman is a postdoctoral research associate at the University of Pittsburgh, working with Johanna Moore on text generation and text analysis for another tutoring system project.

S. Stefan Brandle finished his thesis only a month ago. He has decided to stay in Chicago and write papers with us this next year. He is trying to choose between a Visiting Assistant Professorship at Illinois Institute of Technology and one at Wheaton College.

Michael Glass hopes to finish his thesis in late summer or early fall. He has been discussing a job in Pittsburgh with Lori Levin.

List of Papers Written by AASERT Students

- Brandle, S. & Evens, M. (1997a). Acknowledgements in tutorial dialogue. *Proceedings of MAICS '97*. Dayton, OH, June, 1997. 13-18.
- Brandle, S. & Evens, M. (1997b). Organizing acknowledgements in tutorial dialogue. *Proceedings of the Cognitive Science Conference*, August, Stanford, CA. 872.
- Brandle, S. & Evens, M. (1988). Categorizing acknowledgements in tutorial dialogue. *Proceedings of CogSci '98*, Madison, WI.
- Freedman, R. (1995). Using pedagogical knowledge to structure text generation in an intelligent tutoring system. *Proceedings of the Midwest Artificial Intelligence and Cognitive Science Conference, MAICS '95*, Carbondale, IL, April. 48-52.
- Freedman, R. (1996a). Using a text planner to model the behavior of human tutors in an ITS, *Proceedings of the 1996 Midwest Artificial Intelligence and Cognitive Science Society Conference*, Bloomington, IN. <http://www.cs.indiana.edu/event/maics96/Proceedings/Freedman/freedman.html>
- Freedman, R. (1996b). Using tutoring patterns to generate more cohesive text in an intelligent tutoring system. *Proceedings of International Conference on Learning Systems (ICLS-96)*, Evanston, IL. 75-82.
- Freedman, R. & Evens, M. (1996). Generating and revising hierarchical multi-turn text plans in an ITS. *Proceedings of the Third International Conference on Intelligent Tutoring Systems (ITS '96)*, Montreal, Canada. 632-640.
- Freedman, R. (1997). Degrees of mixed-initiative interaction in an intelligent tutoring system. *Computational Models for Mixed Initiative Interaction, AAAI Spring Symposium*, Stanford University, March 24-26, 44-49.
- Freedman, R. & Evens, M. (1997). The use of multiple knowledge types in an intelligent tutoring system. *Proceedings of the Cognitive Science Conference*, Stanford, CA. 920.
- Freedman, R. 1997. Representing Communicative Action in a Dialogue-Based Intelligent Tutoring System. *Proceedings of AAAI 1997 Fall Symposium: Communicative Action in Humans and Machines*.
- Freedman, R., Zhou, Y., Kim, J.H., Glass, M., & Evens, M. (1998a). SGML-based markup as a step toward improving knowledge acquisition for text generation, *AAAI Spring Symposium on Applying Machine Learning to Discourse Processing*.
- Freedman, R., Zhou, Y., Glass, M., Kim, J.H., & Evens, M. (1998b). Using rule induction to assist in rule construction for a natural-language based intelligent tutoring system. *Proceedings of 20th Annual Cognitive Science Conference*, Madison, WI, August.
- Freedman, R., Brandle, S., Glass, M., Kim, J.H., Zhou, Y., & Evens, M. To appear. Content Planning as the Basis for an Intelligent Tutoring Session. *Proceedings of the Ninth International Workshop on Natural Language Generation (INLG-9)* Niagara on the Lake, Canada.
- Freedman, R., Zhou, Y., Glass, M., Kim, J.H., and Evens, M. To appear. Using Rule Induction to Assist in Rule Construction for a Natural-Language Based Intelligent Tutoring System. *Proceedings of the Twentieth Annual Conference of the Cognitive Science Society*, Madison, WI.

- Glass, M., & Evens, M. (1996). Goals for the Cirsim-Tutor input understander. *Proceedings of MAICS96*, Bloomington, IN. <http://www.cs.indiana.edu/event/maics96/Proceedings/glass.html>
- Glass, M. (1997). Some phenomena handled by the Cirsim-Tutor Version 3 input understander. *Proceedings of the Tenth International Florida Artificial Intelligence Research Symposium*, Daytona Beach, FL. 21-25.
- Hume, G., & Evens, M. (1992). Student modeling and the classification of errors in an intelligent cardiovascular tutoring system, *Proceedings of the Fourth Midwest Artificial Intelligence and Cognitive Science Conference*, Starved Rock, IL. 52-56.
- Hume, G., Evens, M., Rovick, A.A., & Michael, J. (1993). The use of hints as a tutoring tactic. *Proceedings of CogSci '93*. Boulder, CO. 563-568.
- Hume, G., Michael, J.A., Rovick, A.A., & Evens, M. (1995). Controlling active learning: How tutors decide when to generate hints. *Proceedings of FLAIRS '95*. Melbourne Beach, FL. 157-161.
- Hume, G., Michael, J.A., Rovick, A.A., & Evens, M. (1996a). Hinting as a tactic in one-on-one tutoring. *Journal of the Learning Sciences*, Vol. 5, No. 1, 23-47.
- Hume, G., Michael, J.A., Rovick, A.A., & Evens, M. (1996b). Responses and follow up tutorial tactics in an ITS. *Proceedings of the 1996 Florida Artificial Intelligence Research Symposium*. May 20-22, Key West, FL, 168-172.
- Kim, J.H., Freedman, R., & Evens, M. (1998a). Relationship between tutorial goals and sentence structure in a corpus of tutoring transcripts. *Ninth Midwest Artificial Intelligence and Cognitive Science Conference*, Dayton, OH, AAAI Press. 124-131.
- Kim, J.H., Freedman, R., & Evens, M. (1998b). Responding to unexpected student utterances in Cirsim-Tutor v.3: Analysis of transcripts. *FLAIRS-98. Florida Artificial Intelligence Research Symposium*, Sanibel Island, FL. 153-157.
- Sanders, G., Evens, M., Hume, G., Rovick, A.A., & Michael, J.A. (1992). An analysis of how students take the initiative in keyboard-to-keyboard dialogues in a fixed domain. *Proceedings of the Cognitive Science Conference*, Bloomington, August. 1086-1091.